



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
 United States Patent and Trademark Office
 Address: COMMISSIONER FOR PATENTS
 P.O. Box 1450
 Alexandria, Virginia 22313-1450
 www.uspto.gov

| APPLICATION NO. | FILING DATE | FIRST NAMED INVENTOR | ATTORNEY DOCKET NO. | CONFIRMATION NO. |
|---|-------------|----------------------|--------------------------------|------------------|
| 10/694,923 | 10/27/2003 | Michael M. Klock | NVID-062/00US | 5104 |
| 23419 7590 03/27/2007 COOLEY GODWARD KRONISH LLP 3000 EL CAMINO REAL 5 PALO ALTO SQUARE PALO ALTO, CA 94306 | | | EXAMINER WASHBURN, DANIEL C | |
| | | | ART UNIT | PAPER NUMBER |
| | | | 2628 | |
| SHORTENED STATUTORY PERIOD OF RESPONSE | | MAIL DATE | DELIVERY MODE | |
| 3 MONTHS | | 03/27/2007 | PAPER | |

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary

Application No.

10/694,923

Applicant(s)

KLOCK ET AL.

Examiner

Dan Washburn

Art Unit

2628

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 27 October 2003.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-27 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-27 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 27 October 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____.
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____.

DETAILED ACTION

Claim Objections

Claim 12 is objected to because of the following informalities: Claim 12 begins with, "The method of claim 11..." it should read, "The method of *claim 11*..."

Appropriate correction is required.

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

Claims 1-21 and 27 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

Claims 1-21 and 27 read on an abstract idea, as initially defined in claims 1, 11, 21, and 27. For example, claim 1 describes monitoring as a function of time one attributes of a graphics system and selecting a performance level based on the measured attribute, which is not considered transforming an article or physical object to a different state or thing, and is not considered producing a useful, concrete, and tangible result.

Likewise, claims 11, 21, and 27 all describe monitoring and selecting steps, which is not considered transforming an article or physical object to a different state or thing, and is not considered producing a useful, concrete, and tangible result.

Therefore claims 1-21 and 27 are directed solely at an abstract idea.

For claims including such excluded subject matter to be eligible, the claim must be for a practical application of the abstract idea, law of nature, or natural phenomenon.

Art Unit: 2628

Diehr, 450 U.S. at 187, 209 USPQ at 8 ("application of a law of nature or mathematical formula to a known structure or process may well be deserving of patent protection."); Benson, 409 U.S. at 71, 175 USPQ at 676 (rejecting formula claim because it "has no substantial practical application").

To satisfy section 101 requirements, the claim must be for a practical application of the § 101 judicial exception, which can be identified in various ways:

- The claimed invention "transforms" an article or physical object to a different state or thing.
- The claimed invention otherwise produces a useful, concrete and tangible result.

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent; except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1, 2, 4-6, 8-12, 14-16, 18-20, 22, and 25-27 are rejected under 35

U.S.C. 102(e) as being anticipated by Culbert et al. (US 6,820,209).

As to claim 1, Culbert describes a method of operating a graphics system having at least two performance levels, comprising:

monitoring as a function of time at least one attribute of said graphics system that is indicative of a level of graphics activity (column 1 lines 60-67 and column 2 lines 1-10 describes a controller that monitors the activity of components within the controller which require access to local memory, and then decreases a clocking frequency for a memory interface to the local memory when the monitoring indicates that reduced amounts of activity are present. Following such a decrease in the clocking frequency, when increased amounts of activity are detected, the clocking frequency is increased for high performance operation. Column 2 lines 15-67 and column 3 lines 1-11 further describe that the controller is a graphics controller in a graphics display system); and

selecting a performance level for said level of graphical activity to maintain a minimum desired display rate (column 1 lines 60-67 and column 2 lines 1-10 describe that the clocking frequency for the memory interface and the memory is altered according to the measured amount of memory activity, which is considered selecting a performance level, and column 4 lines 42-56 describes that the local memory 122 can, in part, implement a frame buffer that stores the data that is used to produce an image on the display device 120. The local memory is used not only to refresh the display device 120 but also to provide data storage for the processing engines provided within the graphics controller 116 and to receive or exchange data over the system bus 118 under the control of the processing unit 102 or the graphics controller 116. As such, the local memory 122 and the memory bus 124 are heavily used resources of the computer system 100. Further, column 6 lines 19-44 describes that the display interface 210 is required to periodically refresh the display device and thus has a predetermined for

usage of the local memory 202 that stores the data to be displayed. The frequency of the clock for the memory and memory interface is based on the level of graphical activity within the graphics controller, and the memory and memory interface are accessed in order to periodically refresh the display device even when the graphical activity within the system is high, thus, the variable frequency of the clock is considered selecting a performance level for said level of graphical activity to maintain a minimum desired display rate).

Regarding claim 2, Culbert describes the method of claim 1, wherein said monitoring said at least one attribute comprises:

monitoring at least one attribute indicative of utilization of a graphics memory (column 3 lines 55-67 and column 4 lines 1-6).

Concerning claim 4, Culbert describes the method of claim 1, wherein said selecting said performance level comprises:

increasing said performance level in response to detecting an over-utilization condition (column 2 lines 27-43).

With regard to claim 5, Culbert describes the method of claim 1, wherein said selecting said performance level comprises:

decreasing said performance level in response to detecting an under-utilization condition (column 2 lines 27-43).

As to claim 6, Culbert describes the method of claim 1, wherein monitoring said at least one attribute comprises:

monitoring a first attribute indicative of utilization within a graphics processor core clock domain (column 4 lines 35-41 describes that the graphics controller 116 supports, for example, 2D and 3D rendering of graphical images for display on the display device 120. However, the graphics controller 11 can also display text and simple shapes without needing to utilize the sophisticated processing capabilities of the graphics controller with respect to the 2D and 3D images. Column 6 lines 19-35 describes that the 2D and 3D graphics engines are only activated when the graphics controller 200 needs to produce 2D or 3D graphics, in which case access to local memory would be needed. Finally, column 6 lines 56-67 and column 7 lines 1-16 describes that the clock controller can selectively supply a third clock signal (CK3) to the 2D graphics engine and a fourth clock signal (CK4) to the 3D graphics engine when the processing resources of either (or both) engines are required. The clock controller 218 selectively outputting a clock signal to the 2D graphics engine and/or the 3D graphics engine when the 2D and/or 3D graphics engine is required for complex 2D and/or 3D graphical processing is considered monitoring a first attribute (whether or not the 2D and/or 3D graphics engine is required) indicative of utilization within a graphics processor core clock domain (the core clock domain in this case is the clock signals that are selectively sent to the 2D and 3D graphics engines); and

monitoring a second attribute indicative of utilization within a graphics memory clock domain (column 6 lines 56-67 and column 7 lines 1-16 describes that the arbitration unit 208 produces a status signal based on the monitoring of the extent to which local memory 202 access is needed by the resources (which is considered an

attribute). The status signal is supplied to a clock controller 218. Clock controller 218 supplies a first clock (CK1) to the local memory interface 204 and supplies a second clock (CK2) to the local memory 202. The clock controller 218 can thus reduce the first and the second clock when the status signal indicates that the bandwidth load on the local memory 202 is reduced. Figure 3A offers an illustration of the components that make up the clock controller (the memory clock domain in this case is the clock signals (either high frequency or low frequency) that are sent to the local memory and the local memory interface)).

Regarding claim 8, Culbert describes the method of claim 6, wherein said monitoring said second attribute comprises:

monitoring the percentage of clock cycles in a graphics memory for which a memory bandwidth of said graphics memory is inadequate (Figure 5 and column 9 lines 31-53 describes memory interface power management processing 500. If memory utilization is high or low memory utilization hasn't persisted for a predetermined period of time then a high speed clock is output 510 to the local memory interface, but if the low memory utilization has persisted for a predetermined period of time then a low speed clock is output 510 to the local memory interface. This is considered monitoring the percentage of clock cycles in a graphics memory for which a memory bandwidth of said graphics memory is inadequate as the high clock speed is output to the local memory interface when the memory utilization has been high for a least one clock cycle within a predetermined period of time. In other words, when the percentage of clock cycles in a graphics memory for which a memory bandwidth is inadequate (if used with

Art Unit: 2628

the low speed clock) is greater than zero percent (over a predetermined interval prior to the current clock cycle) the high speed clock is used).

Concerning claim 9, Culbert describes the method of claim 1, wherein said selecting said performance level further comprises:

selecting a core clock rate of a graphics processor (column 6 lines 56-67 and column 7 lines 1-16 describe that clock controller 218 can selectively output clock signals to the 2D and 3D graphics engines when one or both of the engines are required for graphical processing. The clock controller 218 selects one clock rate to output to the 2D and 3D graphics engines when the graphics engines are needed (namely the maximum clock rate of the system) and a second clock rate when the graphics engines aren't needed (namely a clock rate of zero), which is considered selecting a core clock rate of a graphics processor).

With regard to claim 10, Culbert describes the method of claim 9, wherein selecting said performance level further comprises:

selecting a clock rate for a graphics memory associated with said graphics processor (column 6 lines 56-67 and column 7 lines 1-16 describe that clock controller 218 selects a clock rate to output to the local memory interface 204 and to the local memory 202 based on the measured extent to which local memory 202 access is needed by processing resources).

As to claim 11, Culbert describes a method of operating a graphics system having at least two performance levels, comprising:

monitoring as a function of time at least one attribute of said graphics system that is indicative of a level of utilization of at least one component of said graphics system for which over-utilization of said at least one component is likely to decrease a display rate of said graphics system below a normal display rate (column 2 lines 27-59 describes a graphics controller that receives status information indicating local memory usage requirements for the local memory of the controller and determines if the usage requirements are below a threshold condition. Column 4 lines 42-56 further describes that the local memory includes a frame buffer that stores the data that is used to produce an image on a display device, and further describes that the local memory is not only used to refresh the display device, but also provides data storage for the processing engines provided within the graphics controller is used when exchanging data over I/O system busses. Thus, the graphics controller monitors the utilization of the local memory, and over-utilization of the local memory is likely to decrease a display rate of said graphics system below a normal display rate); and

selecting a performance level for which said level of utilization is not greater than an over-utilization threshold (column 2 lines 27-43 describes that the graphics processor receives status information indicating local memory usage requirements, determines whether the usage requirements are below a threshold condition, operates the graphics controller to interact with the local memory in accordance with a regular frequency clock when the usage requirements exceed the threshold condition or in accordance with a reduced frequency clock when the local memory usage requirements are below the

threshold condition, which is considered selecting a performance level for which said level of utilization is not greater than an over-utilization threshold).

Regarding claim 12, Culbert describes the method of claim 11, wherein said monitoring said at least one attribute comprises:

monitoring at least one attribute indicative of utilization of a graphics memory (column 2 lines 27-43).

Concerning claim 14, Culbert describes the method of claim 11, wherein said selecting said performance level comprises:

increasing said performance level in response to detecting that said level of utilization is greater than said over-utilization threshold (column 2 lines 27-43).

With regard to claim 15, Culbert describes the method of claim 11, wherein said selecting said performance level comprises:

decreasing said performance level in response to detecting said level of utilization being below an under-utilization threshold (column 2 lines 27-43).

As to claim 16, Culbert describes the method of claim 11, wherein monitoring said at least one attribute comprises:

monitoring a first attribute indicative of utilization within a graphics processor core clock domain (see the corresponding limitation in the rejection of claim 6, as it is identical in scope); and

monitoring a second attribute indicative of utilization within a graphics memory clock domain (see the corresponding limitation in the rejection of claim 6, as it is identical in scope).

Regarding claim 18, Culbert describes the method of claim 16, wherein said monitoring said second attribute comprises:

monitoring the percentage of clock cycles in a graphics memory for which a memory bandwidth of said graphics memory is inadequate (see the rejection of claim 8, as it is identical in scope).

Concerning claim 19, Culbert describes the method of claim 11, wherein said selecting a performance level further comprises:

selecting a core clock rate of a graphics processor (see the rejection of claim 9, as it is identical in scope).

With regard to claim 20, Culbert describes the method of claim 19, wherein selecting said performance level further comprises:

selecting a clock rate for a graphics memory associated with said graphics processor (see the rejection of claim 10, as it is identical in scope).

As to claim 22, Culbert describes a graphics system having at least two performance levels, comprising:

a performance level controller configured to monitor, as function of time, at least one attribute of said graphics system indicative of a level of graphics activity and to select a performance level sufficient to provide a desired minimum display rate (column 5 lines 47-67 and column 6 lines 1-18 describes graphics controller 200 and column 6 lines 19-43 describes arbitration unit 208, which not only arbitrates which of the resources within the graphics controller gain access to the local memory 202 but also monitors the extent to which local memory 202 access is needed by the resources

within the graphics controller 200. Column 4 lines 42-56 describes that the local memory includes a frame buffer that stores the data that is used to produce images on a display device and further describes that the local memory is used not only to refresh the display device, but also to provide data storage for the processing engines provided within the graphics controller. Thus, the arbitration unit monitors the level of graphics activity of the local memory in order to select a performance level (column 6 lines 56-67 and column 7 lines 1-16) sufficient to provide a minimum desired display rate).

Regarding claim 25, Culbert describes a graphics system, comprising:

a graphics processor having at least two performance levels, each performance level having an associated graphics processor core clock rate and a memory clock rate (Figures 1 and 2, column 5 lines 46-67, and column 6 lines 1-43 describe graphics controller 200 (graphics controller 116), which is considered a graphics processor. The graphics controller 200 operates at two performance levels, based on the activity of the local memory (column 3 lines 55-67 and column 4 lines 1-7), and each performance level has an associated graphics processor core clock rate (column 6 lines 56-67 and column 7 lines 1-16 describes that a third clock signal (CK3) and a fourth clock signal (CK4) are selectively sent to the 2D graphics engine and the 3D graphics engine, respectively, when the 2D or 3D graphics engines are required for graphical processing. At the lower performance level the 2D and 3D graphics engines may not be required and thus no clock signal is sent to either engine, which is considered a first clock rate, but at the higher performance level the 2D and/or 3D graphics engine may be required for complex graphical processing and thus the 2D and/or 3D graphics engine is supplied

with a clock signal from the clock controller, which is considered a second clock rate). Each performance level also has an associated memory clock rate (column 6 lines 56-67 and column 7 lines 1-16 describes that the arbitration unit measures the amount of graphical activity in local memory 202 and supplies a signal to clock controller 218. The clock controller 218 then either outputs a lower speed clock signal or a higher speed clock signal to the local memory and the local memory interface, based on the signal received from the arbitration unit));

a graphics memory coupled to said graphics processor by a graphics bus and operable at said memory clock rate (column 5 lines 46-56 describes that the graphics controller 200 can be used with local memory 202. The local memory 202 can either be internal to the graphics controller chip or separate and coupled to the graphics controller chip (see Figures 1 and 2, the local memory is coupled to the graphics controller chip via bus 124 in Figure 1 and via bus 206 in Figure 2));

a performance level controller, said performance level controller configured to monitor, as function of time, at least one attribute of said graphics system indicative of a level of utilization of at least one component of said graphics system for which over-utilization of said component decreases a display rate (column 6 lines 19-67 and column 7 lines 1-16 describes arbitration unit 208, which is considered a performance level controller. Arbitration unit 208 monitors, as a function of time, the extent to which local memory 202 access is needed by the resources within the graphics controller 200. The arbitration unit 208 sends a status signal to the clock controller 218 that indicates the level of activity of the local memory and thus the required clock speed for the local

memory and the local memory interface. The local memory is used to periodically refresh the display device, which means that if the local memory is supplied with a clock speed that can't keep up with all the memory requests then the display rate of the associated display device will decrease);

said performance level controller configured to increase said performance level to avoid over-utilization of said at least component (column 6 lines 56-67 and column 7 lines 1-16 describes that the arbitration unit 208 produces a status signal based on the monitoring of the extent to which local memory 202 access is needed by the resources. If the bandwidth load on the local memory is high then the arbitration unit switches the clock speed over to the higher clock rate (column 2 lines 15-43));

said performance level controller configured to decrease said performance level from a high performance level to a lower performance level to avoid under-utilization of said at least one component (column 6 lines 56-67 and column 7 lines 1-16 describes that the arbitration unit 208 produces a status signal based on the monitoring of the extent to which local memory access is needed by the resources. If the bandwidth load on the local memory is low then the arbitration unit switches the clock speed over to the lower clock rate (column 2 lines 15-43)).

Concerning claim 26, Culbert describes a graphics system having at least two performance levels, comprising:

means for monitoring utilization of at least one component of said graphics system for which over-utilization of said at least one component decreases a display rate (column 6 lines 19-43 describes that the arbitration unit 208 monitors the extent to

which local memory 202 access is needed by the resources within the graphics controller 200); and

means for selecting a performance level sufficient to provide a minimum display rate (column 6 lines 56-67 and column 7 lines 1-5 describes that the arbitration unit sends a signal to the clock controller 218 indicating whether the lower clock rate or the higher clock rate is required for the local memory and the local memory interface based on the detected load on the local memory).

Concerning claim 27, Culbert describes a method of operating a graphics system having at least two performance levels, comprising:

monitoring as a function of time at least one attribute of said graphics system that is indicative of a level of graphics activity (column 6 lines 19-43 describes that the arbitration unit 208 monitors the extent to which local memory 202 access is needed by the resources within the graphics controller 200); and

selecting a performance level for said level of graphical activity to maintain at least one performance criterion within an acceptable range (column 6 lines 56-67 and column 7 lines 1-5 describes that the arbitration unit sends a status signal to the clock controller based on the extent to which local memory access is needed by the resources. The clock controller then increases, decreases, or doesn't change the rate of the clock output to the local memory and local memory interface based on the status signal from the arbitration unit. Setting the clock frequency of the local memory and the local memory interface based on the measured activity of the local memory is considered selecting a performance level for said level of graphical activity to maintain

at least one performance criterion within an acceptable range, where in this case the performance criterion is how quickly the components of the graphics controller can gain access to the local memory).

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 3, 13, 21, 23, and 24 are rejected under 35 U.S.C. 103(a) as being unpatentable over Culbert et al. (US 6,820,209) in view of Williams et al. (US 6,397,343).

As to claims 3 and 13 Culbert doesn't describe a method wherein said monitoring said at least one attribute comprises: monitoring at least one attribute indicative of utilization of a graphics pipeline.

However, Williams describes a method and system that includes a device for dynamic graphics subsystem clock adjustment within a computer system having a CPU and a dedicated graphics subsystem. A system interface is coupled to the graphics subsystem to allow a controller to determine the graphics processing load placed on the graphics subsystem (column 4 lines 11-31 and column 6 lines 33-50). Williams further describes that monitoring said at least one attribute (in this case the graphics processing load placed on the graphics subsystem) comprises: monitoring at least one attribute indicative of utilization of a graphics pipeline (column 6 lines 51-67, column 7

lines 1-20, and column 8 lines 11-22 describes that the device 100 determines the graphics subsystem load by monitoring the processing activity of graphics subsystem 200 via the pipeline control 206 (e.g., by snooping graphics commands and data flowing through a graphics pipeline to determine the activity level of the graphics pipeline)). It would have been obvious to one of ordinary skill in the art at the time of the invention to include in Culbert the system and method of determining the processing load placed on a graphics subsystem by monitoring the level of activity of a graphics pipeline, as taught by Williams, in order to improve the efficiency of the 2D graphics engine and the 3D graphics engine (disclosed in Culbert) by monitoring the pipeline activity within these graphics engines and determining a required clock rate to be passed to these engines based on the monitored activity. The advantage of passing the graphics engines a variable clock rate, rather than simply passing them no clock rate or the maximum clock rate, is that the system can further reduce its power consumption by only passing the graphics engines the maximum clock rate when the maximum clock rate is required, and otherwise passing the graphics engines a reduced clock rate when the reduced clock rate will still allow the graphics engines to process graphical data in the required amount of time.

Regarding claim 21, Culbert describes a method of operating a graphics system having a high performance level for processing complex three-dimensional graphical images and at least one lower power, lower performance level for processing less complex graphical images, the method comprising:

monitoring as a function of time attributes of a graphics memory of said graphics system that are indicative of a level of utilization of said graphics system (column 2 lines 27-43 describes a controller that monitors the local memory usage requirements for the local memory of the controller);

in response to detecting a level of utilization greater than an over-utilization threshold for which a display rate of the graphics system is likely to be significantly decreased below a normal display rate, selecting a higher performance level (column 2 lines 27-43 describes that the graphics controller interacts with the local memory using a regular frequency clock when the local memory usage requirements exceed a predetermined threshold); and

in response to detecting a level of utilization below an under-utilization threshold, selecting a lower performance level to reduce power required by the graphics system (column 2 lines 27-43 describes that if the local memory usage requirements are below the threshold condition then a reduced frequency clock is used by the controller, which significantly reduces the power required by the graphics system).

Culbert doesn't describe that the graphics system monitors attributes of a graphics pipeline that are indicative of a level of utilization of said graphics system.

However, Williams describes a graphics system that monitors attributes of a graphics pipeline that are indicative of a level of utilization of said graphics system. See the rejection of claims 3 and 13 for a discussion of how Williams reads on this limitation and for motivation to combine Williams with Culbert.

Concerning claim 23, Culbert describes the graphics system of claim 22, wherein said performance level controller monitors a graphics memory interface (column 6 lines 56-65 describes that arbitration unit 208 produces a status signal based on the monitoring of the extent to which local memory 202 access is needed by the resources of the graphics controller 200. The status signal is supplied to the clock controller 218 and the clock controller supplies a first clock signal (CK1) to the local memory interface 204 and supplies a second clock signal (CK2) to the local memory 202. The arbitration unit 208 is therefore considered a performance level controller that monitors a graphics memory interface).

Culbert doesn't describe a system wherein the performance level controller monitors a graphics pipeline.

However, Williams describes a system wherein a performance level controller monitors a graphics pipeline. See the rejection of claims 3 and 13 for a detailed explanation of how Williams reads on this limitation and for motivation to combine Williams with Culbert.

With regard to claim 24, Culbert describes the graphics system of claim 22, wherein said at least two performance levels comprise:

a low power graphics level having a first core clock rate and a first memory clock rate (column 6 lines 19-67 and column 7 lines 1-16 describes that the arbitration unit 208 monitors the usage of the local memory and determines if the 2D and 3D graphics engines should be activated (i.e., the graphics engines should be supplied with the maximum frequency clock signal from the clock controller). If the arbitration unit

determines that the 2D and 3D engines aren't required for the current graphical processing then the arbitration unit sends a signal to the clock controller indicating that no clock should be supplied to the 2D and 3D units. Further, if the arbitration unit determines that the activity at the local memory is low enough that a reduced frequency clock signal is sufficient to carry out all required processing then a reduced frequency clock signal is sent to the local memory and local memory interface. The graphics controller operating in this state is considered a graphics system operating at a low power graphics level having a first core clock rate (in this case a rate of zero) and a first memory clock rate (in this case the reduced clocking frequency)); and

a high performance three-dimensional graphics level having a second core clock rate and a second memory clock rate (column 6 lines 19-67 and column 7 lines 1-16 describes that the arbitration unit 208 monitors the usage of the local memory and determines if the 2D and 3D graphics engines should be activated. If the arbitration unit determines that the 3D graphics engine is required for complex 3D graphics processing then the maximum clock rate of the system is supplied to the 3D graphics engine. Further, if the arbitration unit determines that the activity of the local memory is high enough that the maximum frequency clock signal is required to carry out all required processing then the maximum frequency clock signal is sent to the local memory and the local memory interface. The graphics controller operating in this state is considered a graphics system operating at a high performance three-dimensional graphics level having a second core clock rate (in this case the maximum frequency clock rate of the

system) and a second memory clock rate (in this case the maximum frequency clock rate of the system)).

Culbert doesn't describe that the low power graphics level is a low power three-dimensional graphics level.

However, Williams describes a low power three-dimensional graphics level for a system with dynamic clock frequency adjustment (column 11 lines 1-55 and Figures 4 and 5 describe a 3D graphics computer system. The system includes a clock adjustment device 100 that is coupled to each of the components 410-416 via a clock bus 405 and a pipeline control 406. The clock adjustment device 100 is able to determine the activity of the 3D graphics pipeline and adjust the clock signal accordingly (column 7 lines 4-19 and column 8 lines 11-22). It would have been obvious to one of ordinary skill in the art at the time of the invention to include in Culbert a low power three-dimensional graphics level, as taught by Williams, in order to improve the efficiency of the 3D graphics engine (disclosed in Culbert) by monitoring the pipeline activity within the graphics engine and determining a required clock rate to be passed to the engine based on the monitored activity. The advantage of passing the graphics engine a variable clock rate, rather than simply passing it no clock rate or the maximum clock rate, is that the system can further reduce its power consumption by only passing the graphics engine the maximum clock rate when the maximum clock rate is required, and otherwise passing the graphics engine a reduced clock rate when the reduced clock rate will still allow the graphics engine to process graphical data in the required amount of time.

Claims 7 and 17 are rejected under 35 U.S.C. 103(a) as being unpatentable over Culbert et al. (US 6,820,209) in view of Bose et al. (US 7,076,681).

As to claims 7 and 17, Culbert doesn't describe the method of claim 6, wherein monitoring said first attribute comprises: monitoring the percentage of clock cycles in a graphics pipeline for which at least one stage is held up waiting for the results of another stage.

However, Bose describes an integrated circuit such as a scalar processor. Circuit components or units are clocked by and synchronized to a common system clock. A local clock generator in each clocked unit combines the common system clock and stall status from one or more other units to adjust the register clock frequency up or down (column 3 lines 52-62). Bose further describes (column 6 lines 36-60 and Figure 2) an example that includes an instruction unit (I-unit) and an execution unit (E-unit), the I-unit and E-unit include activity monitoring and clock control logic 126, 128, respectively, which monitor unit activity level. When the E-unit 124 senses a stall condition it asserts a stall bit 130, which is used to adjust down the clock speed of the I-unit clock (to throttle down the I-unit and effectively reduce the instruction rate to the E-unit). Depending on the granularity of the control the E-unit activity status or stall bit 130 can adjust its own clock within the E-unit. When the E-unit stall ends the I-unit clock is throttled back up to its normal clock rate. Similarly, when the I-unit experiences a stall condition it sends an I-pipe empty bit to the E-unit so the E-unit can adjust down its clock to conserve power. The activity monitoring and clock control logic 126, 128 contained within each stage of the scalar processor is considered to monitor the

percentage of clock cycles in a graphics pipeline for which at least one stage (e.g., the E-unit) is held up waiting for the result of another stage (e.g., the I-unit). In this case when the percentage of clock cycles for which the E-unit is held up waiting for the results of the I-unit is greater than zero the activity monitoring and clock control logic within the I-unit notifies the E-unit stage to reduce its clock speed in order to conserve power until the stalling I-unit overcomes the stall. It would have been obvious to one of ordinary skill in the art at the time of the invention to include in Culbert the system and method of monitoring the percentage of clock cycles in a graphics pipeline for which at least one stage is held up waiting for the results of another stage, as taught by Bose, in order to improve the efficiency of the 2D graphics engine and the 3D graphics engine (disclosed in Culbert) by monitoring the pipeline activity within these graphics engines and determining a clock rate for each stage within the pipeline. The advantage of assigning a clock rate to each stage within the pipeline, rather than simply passing them a single clock rate, is that the system can further reduce its power consumption by adjusting the clock rate of stages that are waiting for the results of another stage in order to conserve power at those stage until the stalling stage produces its result.

Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure. Zdravkovic (US 6,715,089) and Mizuyabu et al. (US 7,114,086) describe systems that monitor the number of pending instructions in an instruction buffer and adjust the power consumption of the system based on the number of pending instructions, Sinclair et al. (US 6,848,058) describes a power consumption reduction

circuit that monitors memory access activity in order to selectively reduce power consumption, Cui et al. (US 7,149,909) describes a power management circuit adapted to monitor the idleness of the graphics core and reduce a frequency level of the render clock if the idleness exceeds a predetermined percentage of time, Menezes et al. (US 6,845,456) describes a CPU with multiple performance states that periodically obtains utilization information and adjusts the performance state according to the utilization, Williams (US 5,774,704) describes an apparatus with dynamic CPU clock adjustment, and Kawata (US 6,076,171) describes an apparatus with a CPU-load-based clock frequency.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Dan Washburn whose telephone number is (571) 272-5551. The examiner can normally be reached on Monday through Friday 8:30 a.m. to 5:00 p.m..

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Ulka Chauhan can be reached on (571) 272-7782. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.


Art Unit: 2628

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

DW

DW

3/19/07



KEE M. TUNG
SUPERVISORY PATENT EXAMINER